Example Cost Functions – SELECT

algorithm	condition	# blocks read	
SL – linear search (brute force)	equality, key	b/2	on average
	equality, not key range	b	
SB – binary search	equality, key	$\log_2 b$	
	equality, not key	$\log_2 b + \lceil s/bfr \rceil$	
	range	$\log_2 b + b/2$	for range searches, assume roughly half the file records (and thus half the blocks if the file is ordered on the field involved in the condition) will satisfy the condition – can be very inaccurate in specific cases, but reasonably correct on average (can use better estimate if relevant DB stats are available)
SH – hash file	equality, key	1 or 2	
SP – primary index	equality, key	x + 1	
	range	x + b/2	
SC – clustering index	equality	$x + \lceil s/bfr \rceil$	
	range	x + b/2	
SS – secondary index	equality, key	x + 1	
	equality, not key	x + s	
	range	$x + b_1/2 + r/2$	

Number of blocks in result: s/bfr

x = # levels in index, s = selection cardinality (# matches)

Example Cost Functions – JOIN ($R \bowtie_{R.A=S.B} S$)

name	implementation	# blocks read
JNL – nested-loop join	read in n-2 blocks of R at a time read all blocks of S for each n-2 blocks of R, one at a time use one block of memory for assembling result	$b_R + \lceil b_R/(n-2) \rceil b_S$
JSL – single-loop join	read in a block of R, find all matches of S using the index	secondary index: $b_R + \|R\| (x_B+s_B)$
		clustering index: $b_R + \|R\| (x_B+\lceil s_B/bfr_S \rceil)$
		primary index: $b_R + \|R\| (x_B+1)$
		hash key: $b_R + \|R\|$
JSM – sort-merge join	sort files (blocks read and written)	external sorting: $2b + 2b \log_{min(n-1,b/n)} (b/n) \approx 2b \log_2 b$
	merge	$b_R + b_S$

Number of blocks in result: (js |R| |S| / bfr_{RS})

x_B = # levels in index on field B, js = join selectivity, n = # memory blocks to use

Evaluating Query Cost

What information do we need for evaluating cost functions?

- number of records (r), number of blocks (b)
- blocking factor (bfr)
  - can be calculated directly or estimated from b, r
  - bfr_{RS} can be calculated explicitly if column sizes are known, or estimated as $\lceil 1/(b_R/r_R+b_S/r_S) \rceil = \lceil bfr_R bfr_S/(bfr_R+bfr_S) \rceil$
- physical file organization, available indexes, number of levels (x) of multilevel indexes, number of first-level index blocks (b_1)
- number of distinct values (d)
- attribute selectivity (sl), selection cardinality (s) for each attribute
  - can be computed from d and r (requires assumption of uniform distribution or knowledge of distribution if non-key)

Obtaining the Necessary Information

DBMS stores this information.

- frequently-changed values may not be kept completely up-to-date

DB Stats

table	column	# distinct values (d)	low value	high value
PROJECT	Plocation	200	1	200
	Pnumber	2000	1	2000
	Dnum	50	1	50
DEPARTMENT	Dnumber	50	1	50
	Mgr_ssn	50	1	50
EMPLOYEE	Ssn	10000	1	10000
	Dno	50	1	50
	Salary	500	1	500
	Bdate		1945-01-01	1989-12-31

table	# records (r)	# blocks (b)
PROJECT	2000	100
DEPARTMENT	50	5
EMPLOYEE	10000	2000

index	# levels (x)	# level 1 blocks (b_1)
PROJ_PLOC	2	4
EMP_SSN	2	50
EMP_SAL	2	50

Cost of an Execution Plan – Plan 1

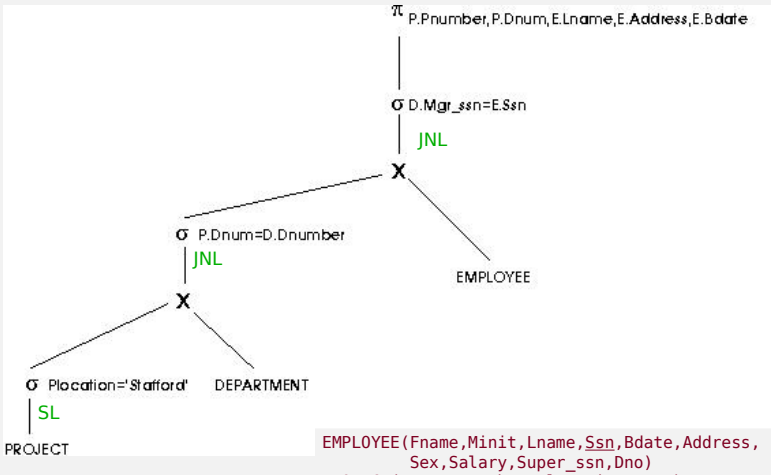


```
EMPLOYEE(Fname,Minit,Lname,Ssn,Bdate,Address,
         Sex,Salary,Super_ssn,Dno)
PROJECT(Pname,Pnumber,Plocation,Dnum)
WORKS_ON(Essn,Pno,Hours)
```


Cost of Plan 1

assuming no pipelining

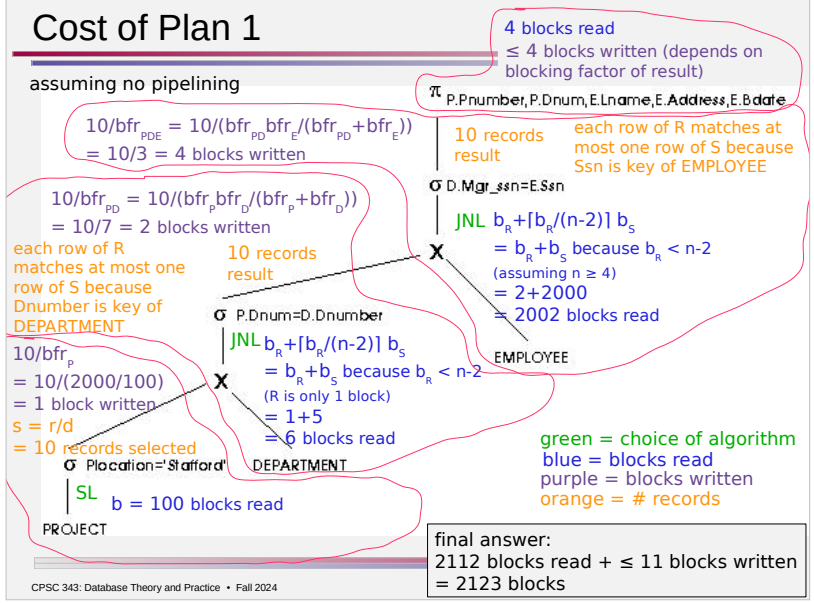

final answer:
2112 blocks read + ≤ 11 blocks written
= 2123 blocks

Cost of Plan 1

with pipelining – differences in **bold**

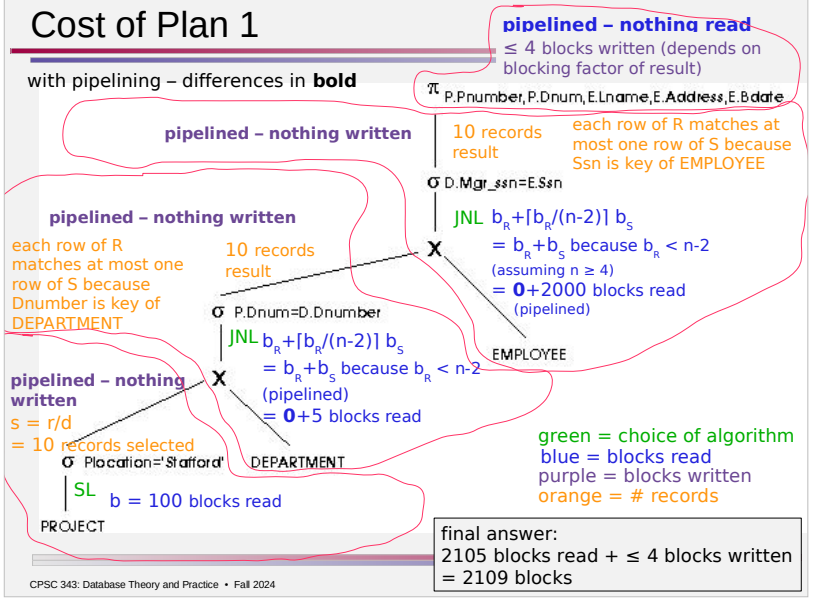

final answer:
2105 blocks read + ≤ 4 blocks written
= 2109 blocks

Cost of an Execution Plan – Plan 2

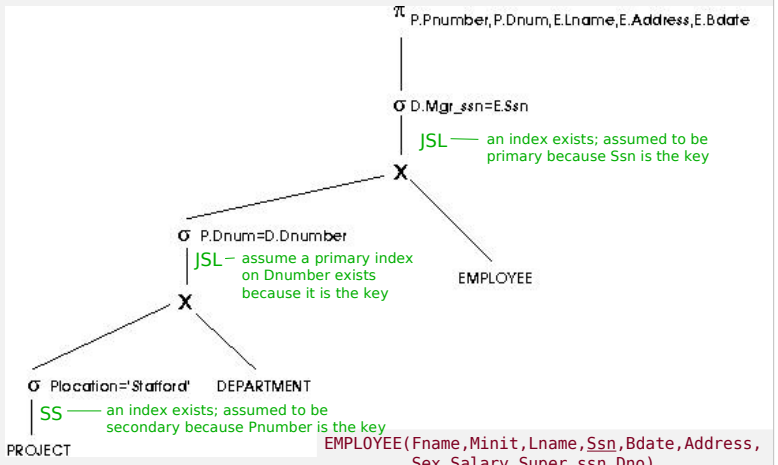


```
EMPLOYEE(Fname,Minit,Lname,Ssn,Bdate,Address,
         Sex,Salary,Super_ssn,Dno)
PROJECT(Pname,Pnumber,Plocation,Dnum)
WORKS_ON(Essn,Pno,Hours)
```


Cost of Plan 2

assuming no pipelining

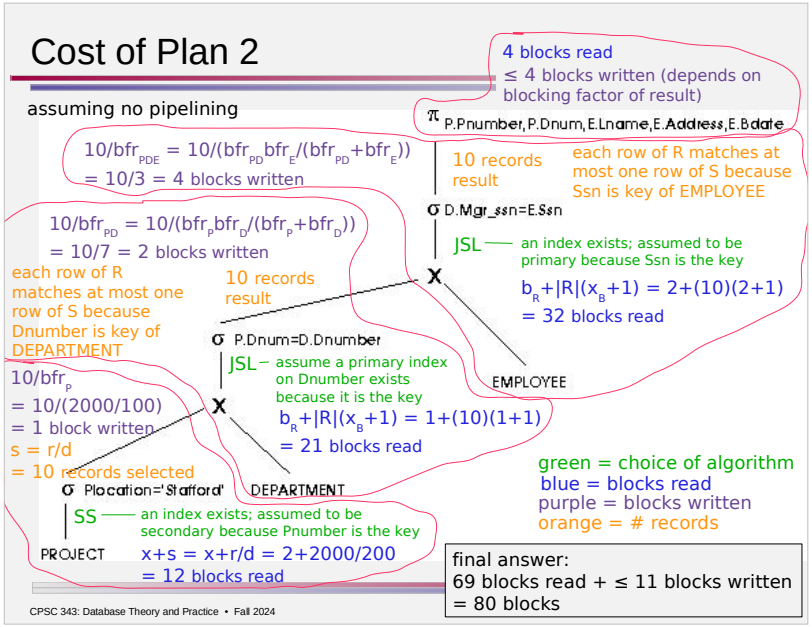

Cost of Plan 2

with pipelining – differences in **bold**